arXiv:1401.6276v1 [stat.ML] 24 Jan 2014

The EM algorithm and the Laplace Approximation

Niko Brümmer
AGNITIO Research South Africa

January 27, 2014

The Laplace approximation calls for the computation of second derivatives at the likelihood maximum. When the maximum is found by the EM algorithm, there is a convenient way to compute these derivatives. The likelihood gradient can be obtained from the EM-auxiliary, while the Hessian can be obtained from this gradient with the Pearlmutter trick.

1 The Laplace approximation

Let X denote the observed data, H some hidden variables and Θ the model parameters. We assume the *joint distribution*:

$$P(X, H, \Theta) = P(X|H, \Theta)P(H|\Theta)P(\Theta) \tag{1}$$

is easy to work with, while the *marginal distribution*:[1]

$$P(X, \Theta) = \int P(X, H, \Theta)\, dH \tag{2}$$

has a more complex form. The Laplace approximation[2] calls for finding the *mode* and *Hessian*, w.r.t. Θ:

$$\hat{\Theta} = \operatorname{argmax} P(X, \Theta), \qquad \text{and} \qquad \mathbf{\Lambda} = \nabla^2 \log P(X, \hat{\Theta}) \tag{3}$$

The approximation is:

$$P(\Theta|X) \approx \mathcal{N}(\Theta|\hat{\Theta}, -\mathbf{\Lambda}^{-1}) \tag{4}$$

[1] All integrals are definite integrals, with fixed boundaries. If H is discrete, the integral can be replaced by summation.

[2] See: Christopher M. Bishop, *Pattern Recognition and Machine Learning* (Information Science and Statistics), Springer, 2007; David J. C. MacKay, *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, 2003.

2 EM-algorithm

If we are using the EM-algorithm for finding the maximum, $\hat{\Theta}$, then the EM-auxiliary provides a convenient route to the Hessian.

2.1 The EM auxiliary

Let Θ' be any valid parameter value satisfying $\int P(H|X,\Theta')\,dH = 1$. We construct the EM auxiliary as follows:

$$\begin{aligned}
\log P(X,\Theta) &= \int P(H|X,\Theta') \log P(X,\Theta)\,dH \\
&= \int P(H|X,\Theta') \log P(X|\Theta)\,dH + \log P(\Theta) \\
&= \int P(H|X,\Theta') \log \frac{P(X,H|\Theta)}{P(H|X,\Theta)}\,dH + \log P(\Theta) \\
&= \int P(H|X,\Theta') \log \left[\frac{P(X,H|\Theta)}{P(H|X,\Theta)} \times \frac{P(H|X,\Theta')}{P(H|X,\Theta')}\right] dH + \log P(\Theta) \\
&= A(\Theta',\Theta) + D(\Theta',\Theta)
\end{aligned} \tag{5}$$

where $A(\Theta',\Theta)$ is the *EM-auxiliary*:

$$A(\Theta',\Theta) = \int P(H|X,\Theta') \log \frac{P(X,H|\Theta)}{P(H|X,\Theta')}\,dH + \log P(\Theta) \tag{6}$$

and $D(\Theta',\Theta) \geq 0$ is KL-divergence:

$$D(\Theta',\Theta) = \int P(H|X,\Theta') \log \frac{P(H|X,\Theta')}{P(H|X,\Theta)}\,dH \tag{7}$$

Notice that if we zero the divergence by choosing $\Theta' = \Theta$, then:

$$\log P(X,\Theta) = A(\theta,\theta) \tag{8}$$

2.2 Algorithm

Although this note is not about the algorithm itself, we very briefly summarize it. An iteration of the EM-algorithm proceeds as follows: Start at Θ_1. The E-step effectively maximizes $A(\Theta',\Theta_1)$ w.r.t. Θ' by simply setting $\Theta' = \Theta_1$, which minimizes (and therefore zeros) the divergence.[3] The M-step now

[3]Here we vary Θ', while Θ and therefore $\log P(X,\Theta)$ remain *fixed*. Then decreasing D must increase A.

maximizes $A(\Theta_1, \Theta)$, w.r.t. the *other* parameter, Θ, usually by zeroing partial derivatives. This gives some value Θ_2, such that $A(\Theta_1, \Theta_2) \geq A(\Theta_1, \Theta_1)$. The net effect of both steps is:

$$\log P(X, \Theta_2) = A(\Theta_2, \Theta_2) \geq A(\Theta_1, \Theta_2) \geq A(\Theta_1, \Theta_1) = \log P(X, \Theta_1) \quad (9)$$

3 Derivatives

We find the Hessian of $\log P(X, \Theta)$ in two steps. First we find the gradient, which we then differentiate again using the Pearlmutter trick.

3.1 Gradient

The gradient of $\log P(X, \Theta)$ coincides with the gradient of the auxiliary. We show how this works.

Let θ denote some component of Θ, then, for any value of Θ', we have:

$$\frac{\partial}{\partial \theta} \log P(X, \Theta) = \frac{\partial}{\partial \theta} A(\Theta', \Theta) + \frac{\partial}{\partial \theta} D(\Theta', \Theta) \quad (10)$$

Note: we are differentiating *only* w.r.t. the components of Θ and *not* w.r.t. those of Θ'. The derivative of the divergence is:

$$\frac{\partial}{\partial \theta} D(\Theta', \Theta) = -\int \frac{P(H|X, \Theta')}{P(H|X, \Theta)} \frac{\partial}{\partial \theta} P(H|X, \Theta)\, dH \quad (11)$$

which conveniently vanishes at $\Theta' = \Theta$:

$$\begin{aligned} \left[\frac{\partial}{\partial \theta} D(\Theta', \Theta)\right]_{\Theta'=\Theta} &= -\int \frac{\partial}{\partial \theta} P(H|X, \Theta)\, dH \\ &= -\frac{\partial}{\partial \theta} \int P(H|X, \Theta)\, dH = -\frac{\partial}{\partial \theta} 1 = 0 \end{aligned} \quad (12)$$

Putting this together, we find:

$$\begin{aligned} \frac{\partial}{\partial \theta} \log P(X, \Theta) &= \left[\frac{\partial}{\partial \theta} A(\Theta', \Theta)\right]_{\theta'=\theta} \\ &= \int P(H|X, \Theta) \frac{\partial}{\partial \theta} \log P(X, H, \Theta)\, dH \end{aligned} \quad (13)$$

For exponential family distributions, the RHS is usually more convenient than the LHS, because now the log directly simplifies $P(X, H, \Theta)$. Also note that it is unnecessary to differentiate the posterior $P(H|X, \Theta)$, or any associated entropy or divergence.

3.1.1 Other derivatives

Just for interest, we mention here that there are two other derivatives that also vanish:

$$\left[\frac{\partial}{\partial\theta'}A(\Theta',\Theta)\right]_{\Theta'=\Theta} = \left[\frac{\partial}{\partial\theta'}D(\Theta',\Theta)\right]_{\Theta'=\Theta} = \left[\frac{\partial}{\partial\theta}D(\Theta',\Theta)\right]_{\Theta'=\Theta} = 0 \quad (14)$$

where θ' is any component of Θ'. This is because at $\Theta' = \Theta$, A is maximized w.r.t. Θ', while D is minimized w.r.t. both arguments. Only $\frac{\partial}{\partial\theta}A$ does not vanish here, because it is not necessarily at the maximum w.r.t. Θ.

3.2 Hessian

We first examine the Hessian analytically. We now consider θ_i, θ_j, both components of Θ and differentiate first w.r.t. the one and then the other:

$$\begin{aligned}
&\frac{\partial^2}{\partial\theta_i\partial\theta_j}\log P(X,\Theta) \\
&= \frac{\partial}{\partial\theta_j}\int P(H|X,\Theta)\frac{\partial}{\partial\theta_i}\log P(X,H,\Theta)\,dH \\
&= \int P(H|X,\Theta)\frac{\partial^2}{\partial\theta_i\partial\theta_j}\log P(X,H,\Theta)\,dH \\
&\quad + \int \frac{\partial}{\partial\theta_j}P(H|X,\Theta)\frac{\partial}{\partial\theta_i}\log P(X,H,\Theta)\,dH \\
&= \left[\frac{\partial^2}{\partial\theta_i\partial\theta_j}A(\Theta',\Theta)\right]_{\Theta'=\Theta} + \int \frac{\partial}{\partial\theta_j}P(H|X,\Theta)\frac{\partial}{\partial\theta_i}\log P(X,H,\Theta)\,dH
\end{aligned} \quad (15)$$

This is the Hessian of the auxiliary plus an extra term that can get messy to derive and implement. The Pearlmutter trick gives a convenient alternative:

3.2.1 Pearlmutter trick

Let $\nabla f(\Theta)$, a column vector, denote the gradient of some multivariate function f, evaluated at Θ. Similarly, let $\nabla^2 f(\Theta)$, a square matrix, denote the Hessian. Then the Pearlmutter trick[4] computes the product of the Hessian with an arbitrary column vector, $\mathbf{v}$, as:

$$\nabla^2 f(\hat{\Theta})\mathbf{v} = \left[\frac{\partial}{\partial\alpha}\nabla f(\hat{\Theta}+\alpha\mathbf{v})\right]_{\alpha=0} \quad (16)$$

[4]Barak A. Pearlmutter, "Fast exact multiplication by the Hessian", Neural Computation, vol. 6, pp. 147160, 1994.

When Θ has n components, the trick must be applied n times, to map out the columns of the Hessian by successively choosing $\mathbf{v} = [1, 0, 0, \ldots]$, $\mathbf{v} = [0, 1, 0, \ldots]$ and so on.

For practical implementation, the gradient using (13) could be derived[5] and coded by hand. When that function is available, the differentiation could be done via forward-mode, algorithmic differentiation. If complex arithmetic is available, then that can be done with minimal coding effort via complex-step differentiation.

[5]The M-step should be based on those same derivatives.